arXiv:1512.05913v1 [stat.CO] 18 Dec 2015

A novel Bayesian strategy for the identification of spatially-varying material properties and model validation: an application to static elastography

P.S. Koutsourelakis

Center for Applied Mathematics, Cornell University

E-mail: pk285@cornell.edu

Abstract. The present paper proposes a novel Bayesian, computational strategy in the context of model-based inverse problems in elastostatics. On one hand we attempt to provide probabilistic estimates of the material properties and their spatial variability that account for the various sources of uncertainty. On the other hand we attempt to address the question of model fidelity in relation to the experimental reality and particularly in the context of the material constitutive law adopted. This is especially important in biomedical settings when the inferred material properties will be used to make decisions/diagnoses. We propose an expanded parametrization that enables the quantification of model discrepancies in addition to the constitutive parameters. We propose scalable computational strategies for carrying out inference and learning tasks and demonstrate their effectiveness in numerical examples with noiseless and noisy synthetic data.

1. Introduction

The extensive use of large-scale computational models poses several challenges in parameter identification in the context of system identification or performing predictive simulations. Medical imaging represents such an application which has attracted significant interest in recent years as the correct identification of material properties can reveal various pathologies (Oberai et al. 2009, Ganne-Carrié et al. 2006) as well as quantitatively assess the progress of various treatments.

Ultrasound elasticity imaging (elastography) has gained prominense in the context of performing medical diagnosis due its accuracy and low cost. It is based on ultrasound tracking of pre- and post-compression images to obtain a map of position changes from which deformations can be inferred. The pioneering work of Ophir and coworkers (Ophir et al. 1991) followed by several clinical studies (Garra et al. 1997, Bamber et al. 2002, Hall et al. 2003, Giuseppetti et al. 2005, Itoh et al. 2006, Thomas et al. 2006, Regner et al. 2006, Burnside et al. 2007, Zhi et al. 2007, Parker et al. 2011) have demonstrated that the resulting strain images typically improve the diagnostic accuracy over ultrasound alone.

Broadly speaking, there are two approaches that are utilized for calculating the constitutive parameters. In the direct approach, the equations of equilibrium are interpreted as equations for the material parameters of interest, where the inferred strains and their derivatives appear as coefficients (Barbone et al. 2010, McLaughlin & Renzi 2006, Albocher et al. 2009). While such an approach provides a computationally efficient strategy that does not require solution over the whole domain nor knowledge of the boundary conditions, it has certain drawbacks. More importantly perhaps, it does not use the raw data (i.e. noisy displacements) but transformed versions i.e. strain fields which arise by applying sometimes ad hoc filtering and smoothing operators. While these might be plausible, in general alter the informational content of the data and make difficult the quantification of the effect of observation noise. This is amplified when strain derivatives are computed, although not all such approaches require them e.g. (Park & Maniatty 2009). Furthermore, the smoothing employed can smear regions with sharply varying properties and hinder proper identification. Finally, it is non-trivial to determine appropriate boundary conditions in terms of the material parameters of interest.

The alternative to direct methods, i.e. indirect, or iterative, as they are most commonly referred to, admit an inverse problem formulation where the discrepancy (in various norms, (Gockenbach & Khan 2005, Gockenbach et al. 2008)) between observed and model-predicted displacements is minimized with respect to the material fields of interest (Zhang et al. 2006, Oberai et al. 2004, Fehrenbach et al. 2006, Doyley et al. 2006, Liu et al. 2005, Liew & Pinsky 2005, Banerjee et al. 2009, Gokhale et al. 2008, Arnold et al. 2010, Olson & Throne 2010, Bonnet & Constantinescu 2005). While these approaches utilize directly the raw data, they generally imply an increased computational cost as the forward problem and potentially derivatives have to be

solved/computed several times. This effort is amplified when stochastic/statistical formulations are employed as those arising from the Bayesian paradigm, whose cost is comparable to that of a deterministic global optimization technique (Kaipio & Somersalo 2005).

Bayesian techniques are advocated in this paper due their ability to quantify the effect of various sources of uncertainty to the hypotheses tested or the inferences made. One source of uncertainty is obviously the noise in the data which constitutes probabilistic estimates more rational. This is particularly important when multiple hypotheses are consistent with the data or the level of confidence in the estimates produced needs to be quantified. Another source of uncertainty which is largely unaccounted for, is *model uncertainty* (Higdon et al. 2008). Namely, the parameters, whose values are estimated, are associated with a particular forward model about the behavior of the medium (in our case a system of PDEs consisting of equilibrium and constitutive equations) but one cannot be certain about the validity of the model employed. In general, there will be deviations between the physical reality where measurements are made, and the idealized mathematical/computational description. Especially in the context of medical applications, it is crucially important to account for the model discrepancy or inadequacy in order to infer the right material properties and make accurate diagnoses. ‡ Non-intrusive Bayesian strategies, i.e. those that basically make use of the forward model as a black-box, capture model discrepancy with regression models (e.g. Gaussian processes) which are not easily physically-interpretable and cumbersome or impractical when they depend on a large number of input parameters. (Kennedy & O'Hagan 2001, Higdon et al. 2008). In contrast, our approach is intrusive. This enables us to overcome the aforementioned limitations and allows us to directly infer the stresses/pressure in the context of *elastostatics*.

The rest of the paper is organized as follows. Section 2 is devoted to the presentation of the novel Bayesian framework proposed in the context of elastostatics. Section 2.1 discusses computational aspects related to inference techniques for sampling from the posterior and learning schemes for estimating parameter values. Finally section 3 presents numerical results under static plane stress conditions using noiseless and noisy data with particular emphasis on quantifying model discrepancy.

2. Proposed Methodology

The presentation of the ideas in this paper is centered around solid mechanics, in particular elastostatics, but the framework introduced can be directly transitioned to other continua. We discuss first the formulation of the probabilistic model proposed and in subsection 2.1 the inference and learning tasks associated with this description. We adopt a physically-inspired strategy that focuses on quantifying model discrepancies in the context of the *constitutive equation*. From a deterministic point of view

‡ "I remember my friend Johnny von Neumann used to say, 'with four parameters I can fit an elephant and with five I can make him wiggle his trunk.'" *A meeting with Enrico Fermi*, Nature 427, 297; 2004.

it resembles techniques such as constitutive relation error (CRE) or error in the constitutive equation (ECE) that have been developed for a posteriori error estimation and the solution of overspecified inverse problems (Kohn & Vogelius 1984, Ladéveze & Chouaki 1999, Bonnet & Constantinescu 2005, Deraemaeker et al. 2004, Feissel & Allix 2007). We use the term constitutive equations to refer in general to relations between conjugate thermodynamic variables, i.e. stress and strain in solid mechanics or velocity and pressure in flow through permeable media or flux and temperature in heat diffusion.

In the formulations proposed, the constitutive relation supplements the observables and an augmented state space is used that includes all conjugate variables. As it is demonstrated in the sequence, the addition of these unknown parameters simplifies inference tasks and enables the quantification of model errors. The motivation for such an approach stems from the fact that inverse problems in the context of continuum models consist of:

- a *conservation law* that arises from physical principles that are generally well-founded and trusted. In the case of single-phase flow through a porous medium this amounts to the conservation of mass, in solid mechanics to the conservation of linear momentum. In elastostatics in particular this is written as:

$$\nabla \cdot \tilde{\sigma}(\mathbf{x}) + \mathbf{b}(\mathbf{x}) = \mathbf{0}, \quad \mathbf{x} \in \Omega \tag{1}$$

  where $\tilde{\sigma}(\mathbf{x})$ is the stress tensor, $\mathbf{b}$ the body force and Ω the problem domain. Discretized versions of the aforementioned PDE are employed which naturally introduce *discretization* error. This is generally well-studied in the context of linear problems and several a priori (and a posteriori) error estimates are available. In this work we will ignore the discretization error in Equation (1) which corresponds to the verification stage and focus on the validation and calibration aspects.

- a *constitutive law* that is by-and-large phenomenological and therefore provides the primary source of *model uncertainty*. This is represented by the conductivity tensor in heat diffusion, the permeability tensor in flow through porous media or the elasticity tensor $\mathbf{D}$ in solid mechanics:

$$\boldsymbol{\sigma}(\mathbf{x}) = \mathbf{D}(\mathbf{x})\boldsymbol{\epsilon}(\mathbf{x}), \quad \forall \mathbf{x} \in \Omega \tag{2}$$

  where $\boldsymbol{\sigma}(\mathbf{x})$ is the vector of stresses and $\boldsymbol{\epsilon}(\mathbf{x})$ the vector of strains.

- boundary/initial conditions or observables in general (which might include interior displacements). The available data are contaminated by noise and represent the main source of *observation errors*.

In the Bayesian setting advocated, the goal is to evaluate the *posterior density* for the material parameters (i.e. $\mathbf{D}(\mathbf{x})$) as well as quantitatively assess the validity of the aforementioned constitutive relation (Equation (2)).

The numerical implementation requires discretization of the aforementioned equations. For economy of notation, we consider the simplest perhaps discretization

consiting of a finite element triangulation $\mathcal{T}$ of the problem domain Ω using n_{el} constant-strain/stress elements §. If e denotes the element number, the parameters in the formulation proposed are:

- the stress vectors $\boldsymbol{\sigma}_e,\ e = 1,\ldots,n_{el}$ ($3-$ dimensional under plane stress/strain conditions or 6-dimensional in general three-dimensional problems) which are jointly denoted by $\boldsymbol{\sigma} = [\boldsymbol{\sigma}_1,\ldots,\boldsymbol{\sigma}_{n_{el}}]^T$.
- the global displacement vector $\mathbf{u}$. If $\mathbf{u}_e$ denotes the nodal displacement vector of element e then we represent by $\mathbf{L}_e$ the Boolean matrices that relate local and global displacement vectors i.e. $\mathbf{u}_e = \mathbf{L}_e\mathbf{u}$. We further denote by $\boldsymbol{\epsilon}_e$ the element strain vector which is related to $\mathbf{u}_e$ as $\boldsymbol{\epsilon}_e = \mathbf{B}_e\mathbf{u}_e$ where $\mathbf{B}_e$ is the well-known strain-displacement matrix.
- the local constitutive matrices $\mathbf{D}_e$ that relate stress and strains over element e, i.e. $\boldsymbol{\sigma}_e = \mathbf{D}_e\boldsymbol{\epsilon}_e$. These are assumed constant over each element but they could be assigned different values at the nodes of the mesh or integration points of each element.

We will further assume that *noisy* displacement data (at interior or boundary points) are provided and will be denoted by $\mathbf{u}_Q \in \mathbb{R}^{n_Q}$. It is assumed that the observed nodal displacements are given by $\mathbf{Q}\ \mathbf{u}$ where $\mathbf{Q}$ is an appropriate Boolean matrix (if all displacements are observed at all the nodes then $\mathbf{Q} = \mathbf{I}$). Assuming Gaussian noise with variance ν^2, the *likelihood* of $\mathbf{u}_Q$ given $\mathbf{u}$ is normal and:

$$p(\mathbf{u}_Q \mid \mathbf{u}) \propto \frac{1}{\nu^{n_Q}} \exp\{-\frac{1}{2\nu^2}(\mathbf{u}_Q - \mathbf{Q}\mathbf{u})^T(\mathbf{u}_Q - \mathbf{Q}\mathbf{u})\} \tag{3}$$

The observation noise variance ν^2 can be known or unknown in which case we propose employing a conjugate $inverse-Gamma$ hyperprior with hyperparameters (α_ν, β_ν), i.e:

$$p(\nu^2) \propto (\nu^{-2})^{\alpha_\nu - 1} e^{-\beta_\nu/\nu^2} \tag{4}$$

Naturally more complex models that can capture perhaps the spatial dependence of ν can be employed. In general, non-essential boundary conditions might be available as well, i.e. tractions might be prescribed at part of the boundary $\partial\Omega_N \subset \partial\Omega$ i.e.:

$$\mathbf{n} \cdot \tilde{\sigma}(\mathbf{x}) \mid_{\Omega_N} = \boldsymbol{\tau}(\mathbf{x}), \quad \mathbf{x} \in \partial\Omega \tag{5}$$

Noise in these observations could also be added but we omit this to simplify notation.

In the proposed framework, apart from the aforementionned observations, the *data* or likelihood consist also of model-related equations, i.e. the conservation law (Equation (1)) which in the case of standard Bubnov-Galerkin finite element schemes is enforced weakly as:

$$\int_\Omega \boldsymbol{\epsilon}(\mathbf{w}) \cdot \boldsymbol{\sigma} d\mathbf{x} = \int_\Omega \mathbf{w} \cdot \mathbf{b}\ d\Omega + \int_{\partial\Omega_N} \mathbf{w} \cdot \boldsymbol{\tau}\ d\Gamma \tag{6}$$

§ For more complex elements/discretizations, the ensuing formulations can be readily applied if instead we consider each integration point in the element

where $\boldsymbol{\epsilon}(\mathbf{w})$ denote the strains associated with the the weighting functions $\mathbf{w} \in H_0^1(\Omega)$. It is noted that other discretization schemes such as finite volume or discontinuous Galerkin can also be used to enforce the conservation law. with small alterations. In the triangulation $\mathcal{T}$ adopted for discretizing the solution and the weighting functions $\mathbf{w}$ this reduces to:

$$\hat{\mathbf{B}}^T \boldsymbol{\sigma} = \mathbf{f} \tag{7}$$

where $\mathbf{f}$ is the force vector and:

$$\hat{\mathbf{B}}^T = \sum_{e=1}^{n_{el}} (\mathbf{L}_e)^T \int_{\Omega_e} (\mathbf{B}_e)^T \, d\mathbf{x} = \sum_{e=1}^{n_{el}} V_e (\mathbf{L}_e)^T (\mathbf{B}_e)^T \tag{8}$$

where V_e is the volume of element e.

The second model equation relates to the constitutive law which we propose enforcing for every element probabilistically. If the true constitutive law (which is unknown) is different from the one prescribed in Equation (2), then there will be a discrepancy/error $\mathbf{c}_e$ between the actual stresses $\boldsymbol{\sigma}_e$ and the model-predicted stresses $\mathbf{D}_e \boldsymbol{\epsilon}_e = \mathbf{D}_e \mathbf{B}_e \mathbf{u}_e$:

$$\mathbf{c}_e = \boldsymbol{\sigma}_e - \mathbf{D}_e \mathbf{B}_e \mathbf{u}_e \tag{9}$$

Since $\mathbf{c}_e$ is unknown and in accordance with the Bayesian formulation advocated, we propose a hierarchical prior model where:

$$\begin{aligned} &\mathbf{c}_e \mid \boldsymbol{\sigma}_e, \mathbf{u}_e, \boldsymbol{\Sigma}_e \sim \mathcal{N}(\boldsymbol{\sigma}_e - \mathbf{D}_e \mathbf{B}_e \mathbf{u}_e, \boldsymbol{\Sigma}_e) \\ &or \\ &p(c_e \mid \boldsymbol{\sigma}_e, \mathbf{u}_e, \boldsymbol{\Sigma}_e) \propto \tfrac{1}{|\boldsymbol{\Sigma}_e|^{1/2}} exp\{-\tfrac{1}{2}(\boldsymbol{\sigma}_e - \mathbf{D}_e \mathbf{B}_e \mathbf{u}_e)^T \boldsymbol{\Sigma}_e^{-1} (\boldsymbol{\sigma}_e - \mathbf{D}_e \mathbf{B}_e \mathbf{u}_e)\} \end{aligned} \tag{10}$$

In this work we consider a special form of the covariances $\boldsymbol{\Sigma}_e = \lambda_e^2 \mathbf{I}$. The hyperparameters λ_e^2 express the variability of the constitutive error and their magnitude quantifies the *model discrepancy* over each element e. The inferred values λ_e^2 will reveal elements where the model error is high and refinement/improvement is needed. Note for example that if the elastic properties vary within an element e, the corresponding λ_e^2 will be non-zero even if no noise exists in the data. When different discretization schemes are used which might employ higher-order shape functions, distinct λ_e^2 for each integration point can be introduced. The normal prior for $\mathbf{c}_e$ (Equation (10)) is not the only option and was selected here for computational convenience due to its conjugacy with the other distributions as it will be seen in the sequel. It would certainly be worth-while to investigate alternative prior models.

Since the hyperparameters λ_e^2 are unknown, prior models can be employed as well. In this study we make of a Gaussian Markov Random Field (GMRF, (Besag et al. 1991, Besag & Green 1993)) prior which accounts for the fact that the magnitude of the model errors are expected to be spatially correlated. In particular, and since $\lambda_e^2 \geq 0$ we define the prior implicitly through the vector $\mathbf{Z} = \{z_e\}_{e=1}^{n_{el}}$ where $z_e = \log \lambda_e^2$:

$$p(\boldsymbol{\Lambda}) \propto \exp\{-\frac{1}{2} \mathbf{Z}^T \mathbf{W} \mathbf{Z}\} \tag{11}$$

The precision matrix is given by $\mathbf{W} = \frac{1}{\sigma_z^2}\mathbf{H}$ where σ_z^2 is a scale parameter and $\mathbf{H} = [H_{e_1,e_2}]$:

$$H_{e_1,e_2} = \begin{cases} \sum_{e_2=1}^{n_{el}} h_{e_1,e_2} & \text{if } e_1 = e_2 \\ -h_{e_1,e_2} & \text{otrherwise} \end{cases} \tag{12}$$

where $h_{e_1,e_2} > 0$ is a mesure of proximity between elements e_1 and e_2. In this work this was defined with respect to the distance d_{e_1,e_2} between the element centroids as $h_{e_1,e_2} = e^{-d_{e_1,e_2}/d_0}$ where d_0 is a correlation-length parameter. The aforementioned model represents an intrinsic autoregressive prior (Kunsch 1987, Besag & Kooperberg 1995), which is an improper distribution (since $\mathbf{W}$ is semi-positive definite) that has been extensively used in spatial statistics. In particular, since $\sum_{e_2} W_{e_1,e_2} = 0 \ \forall e_1$, it can be easily established that $p(\mathbf{\Lambda})$ penalizes the "jumps" in $\mathbf{Z}$ at neighboring elements, i.e.:

$$p(\mathbf{\Lambda}) \propto \exp\{ \sum_{e_1<e_2} W_{e_1,e_2}(z_{e_1} - z_{e_2})^2 \} \tag{13}$$

It is noted finally that values for the parameters (d_0, σ_z^2) are provided in the numerical results section.

The combination of Equations (3), (4), (7), (10) and (11) leads to the *posterior density* on the model parameters $\mathbf{\Theta} = (\nu^2, \boldsymbol{\sigma}, \{\mathbf{D}_e\}_{e=1}^{n_{el}}, \mathbf{u}, \mathbf{\Lambda} = \{\lambda_e^2\}_{e=1}^{n_{el}})$. In addition to the obervations $\mathbf{u}_Q$, the posterior on $\mathbf{\Theta}$ is *explicitly* conditioned on the model equations i.e. the discretized equation of equilibrium and the constitutive law ‖:

$$\begin{aligned} \pi(\mathbf{\Theta}) = p(\mathbf{\Theta} \mid \mathbf{u}_Q, \mathcal{M}) = \ & p(\mathbf{u}_Q \mid \mathbf{u}, \nu^2) p(\nu^2) \\ & 1_{\{\hat{\mathbf{B}}^T \boldsymbol{\sigma} = \mathbf{f}\}}(\mathbf{\Theta}) \\ & \prod_{e=1}^{n_{el}} p(\mathbf{c}_e \mid \boldsymbol{\sigma}_e, \mathbf{u}_e, \lambda_e^2) \ p(\mathbf{\Lambda}) \\ & p(\mathbf{u}) \end{aligned} \tag{14}$$

The indicator function $1_{\{\hat{\mathbf{B}}^T \boldsymbol{\sigma} = \mathbf{f}\}}(\mathbf{\Theta})$ implies that the support of the distribution includes only stress vectors that satisfy the (discretized) equilibrium equations in Equation (7).

A prior model could also be adopted with repsect to the constitutive parameters $\mathbf{D}_e$. Such priors apart from improving the regularity of posterior are also physcially plausible as one would expect the constitutive properties at neighboring locations to be correlated. Naturally several such models have been proposed in the literature (Kaipio & Somersalo 2005). In this work however this was found unnecessary as the formulation proposed provides a natural correlation between $\mathbf{D}_e$ through the dispacements $\mathbf{u}$ and stresses $\boldsymbol{\sigma}$ which are themselves spatially correlated due to the equilibrium and constitutive equations. This is evident in the conditional posteriors presented in the sequence. In contrast a prior model was adopted for the displacement vector, denoted by $p(\mathbf{u})$ in Equation (14). This can be useful when the observed displacements are sparse or restricted to a portion of the problem domain but its primary utility in the examples contained in section 3 was found to be the regularization of the displacement

‖ this conditioning is denoted by $\mathcal{M}$ in Equation (14)

field in the presence of noise. In particular we adopted an intrinsic autoregressive model as the one employed for $\mathbf{\Lambda}$ in Equation (11):

$$p(\mathbf{u}) \propto \exp\{-\frac{1}{2}\mathbf{u}^T\mathbf{V}\mathbf{u}\} \tag{15}$$

where $\mathbf{V} = \frac{1}{\sigma_u^2}\mathbf{J}$. The matrix $\mathbf{J}$ defined exactly as $\mathbf{H}$ in Equation (12) with proximity between two arbitrary entries u_i , u_j defined with respect to the nodal distance.

It is worth emphasizing that the proposed model and associated posterior contain two sets of additional parameters as compared to tradional Bayesian formulations of the inverse problem: a) the stress vector $\boldsymbol{\sigma}$, and b) the model discrepancy parameters λ_e^2. The introduction of the former enables the quantification of the model discrepancy. Despite the augmented set of parameters, these additional vectors play the role of auxiliary variables that expedite the exploration of the posterior using Gibbs sampling (Higdon 1997) as it will is discussed in subsection 2.1. One can readily obtain, *conditional* posterior densities for all the parameters appearing in $\mathbf{\Theta}$. In particular:

- For ν^2:

$$\nu^{-2} \mid \mathbf{u} \sim Gamma\left(\alpha_\nu + \frac{n_q}{2}, \beta_\nu + \frac{1}{2} \parallel \mathbf{u_Q} - \mathbf{Q}\ \mathbf{u} \parallel^2\right) \tag{16}$$

- For $\mathbf{u}$:

$$\mathbf{u} \mid \nu^2, \boldsymbol{\sigma}, \{\mathbf{D}_e, \lambda_e^2\}_{e=1}^{n_{el}} \sim \mathcal{N}(\mu_u, \mathbf{C}_u) \tag{17}$$

where:

$$\begin{array}{l} \mathbf{C}_u^{-1} = \mathbf{C}^T\mathbf{\Lambda}^{-1}\mathbf{C} + \frac{1}{\nu^2}\mathbf{Q}^T\mathbf{Q} + \mathbf{V} \\ \mu_u = \mathbf{C}_u\left(\mathbf{C}^T\mathbf{\Lambda}^{-1}\boldsymbol{\sigma} + \frac{1}{\nu^2}\mathbf{Q}^T\mathbf{u}_Q\right) \end{array} \tag{18}$$

The aforementioned matrices $\mathbf{C}$ and $\mathbf{\Lambda}$ arise from the model discrepancy terms in Equation (14) as follows:

$$\begin{array}{l} \mathbf{C} = \begin{bmatrix} \mathbf{D}_1\mathbf{B}_1\mathbf{L}_1 \\ \mathbf{D}_2\mathbf{B}_2\mathbf{L}_2 \\ \dots \\ \mathbf{D}_{n_{el}}\mathbf{B}_{n_{el}}\mathbf{L}_{n_{el}} \end{bmatrix} \\ \mathbf{\Lambda} = \begin{bmatrix} \lambda_1^2\mathbf{I} & \mathbf{0} & \dots \mathbf{0} & \\ \mathbf{0} & \lambda_2^2\mathbf{I} & \dots & \dots \\ \mathbf{0} & \mathbf{0} & \dots & \lambda_{n_{el}}^2\mathbf{I} \end{bmatrix} \end{array} \tag{19}$$

- For $\boldsymbol{\sigma}$:

$$\boldsymbol{\sigma} \mid \boldsymbol{u}, \{\mathbf{D}_e, \lambda_e^2\}_{e=1}^{n_{el}} \sim \mathcal{N}(\mu_\sigma, \mathbf{C}_\sigma) \tag{20}$$

where:

$$\begin{array}{l} \mathbf{C}_\sigma = \mathbf{\Lambda} + (\hat{\mathbf{B}}^T\mathbf{\Lambda})^T\left(\hat{\mathbf{B}}^T\mathbf{\Lambda}\hat{\mathbf{B}}\right)^{-1}(\hat{\mathbf{B}}^T\mathbf{\Lambda}) \\ \boldsymbol{\mu}_\sigma = \mathbf{C}\mathbf{u} + (\hat{\mathbf{B}}^T\mathbf{\Lambda})^T\left(\hat{\mathbf{B}}^T\mathbf{\Lambda}\hat{\mathbf{B}}\right)^{-1}(\mathbf{f} - \hat{\mathbf{B}}^T\mathbf{C}\mathbf{u}) \end{array} \tag{21}$$

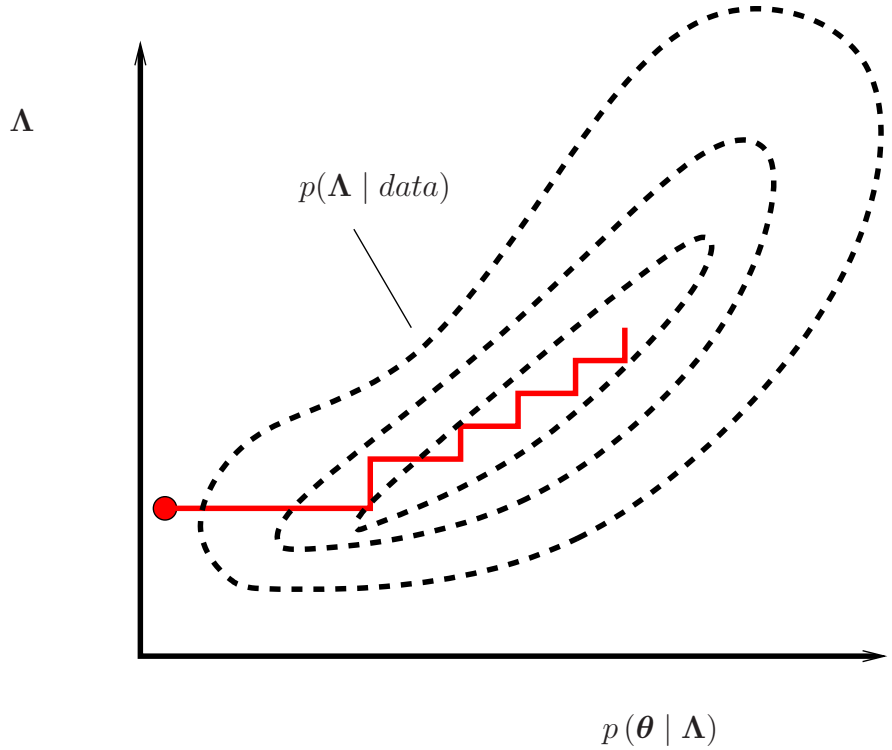

Figure 1. Schematic illustration of the Expectation-Maximization scheme

- For $\mathbf{D}_e$ assuming we are interested in the elastic modulus E_e such that $\mathbf{D}_e = E_e \hat{\mathbf{D}}_e$ (where $\hat{\mathbf{D}}_e$ is known):

$$E_e \mid \boldsymbol{\sigma}_e, \mathbf{u}_e, \lambda_e^2 \sim \mathcal{N}(\mu_E, \sigma_E^2) \tag{22}$$

  where:

$$\begin{aligned} \sigma_E^2 &= \frac{\lambda_e^2}{\|\hat{\mathbf{D}}_e \boldsymbol{\epsilon}_e\|^2} \\ \mu_E &= \frac{\boldsymbol{\epsilon}_e^T \hat{\mathbf{D}}_e^T \boldsymbol{\sigma}_e}{\|\hat{\mathbf{D}}_e \boldsymbol{\epsilon}_e\|^2} \end{aligned} \tag{23}$$

In the following we propose a hybrid scheme based on the Expectation-Maximization algorithm (Dempster et al. 1977) that provides maximum a posteriori (MAP) point estimates for the model discrepancy parameters $\boldsymbol{\Lambda} = \{\lambda_e^2\}$ while fully sampling from the posterior of Equation (14) for the remaining parameters $\boldsymbol{\theta} = (\nu^2, \boldsymbol{\sigma}, \{\mathbf{D}_e\}_{e=1}^{n_{el}}, \mathbf{u})$ (Figure 1).

2.1. Inference and learning

We advocate a scalable procedure for carrying out inference and learning with respect to the posterior $\pi(\boldsymbol{\Theta})$ (Equation (14)) which is common practice in pertinent probabilistic models (Ghahramani 2001). We compute point estimates for the vector $\boldsymbol{\Lambda} = \{\lambda_e^2\}$ which

correspond to maxima $\boldsymbol{\Lambda}^*$ of the *log-posterior*.

$$\begin{aligned} L(\boldsymbol{\Lambda}) = \log p(\boldsymbol{\Lambda} \mid \mathbf{u}_Q, \mathcal{M}) \quad &= \log \int \underbrace{p(\boldsymbol{\Lambda}, \boldsymbol{\theta} \mid \mathbf{u}_Q, \mathcal{M})}_{posterior \;\; Equation(14)} \; d\boldsymbol{\theta} \\ &= \log \int \pi(\boldsymbol{\Lambda}, \boldsymbol{\theta}) \; d\boldsymbol{\theta} \end{aligned} \tag{24}$$

while the remaining parameters $\boldsymbol{\theta} = (\nu^2, \boldsymbol{\sigma}, \{\mathbf{D}_e\}_{e=1}^{n_{el}}, \mathbf{u})$ are sampled from the full posterior $\pi(\boldsymbol{\theta}, \boldsymbol{\Lambda}^*)$.

Maximization of $L(\boldsymbol{\Lambda})$ is more complex than a standard optimization task as it involves integration over the unobserved variables $\boldsymbol{\theta}$. We propose therefore adopting an Expectation-Maximization framework (EM) which is an iterative, robust scheme that is guaranteed to increase the log-posterior at each iteration (Dempster et al. 1977, Ghahramani 2001). It is based on constructing a series of increasing lower bounds of the log-posterior using auxiliary distributions $q(\boldsymbol{\theta})$:

$$\begin{aligned} L(\boldsymbol{\theta}) = \quad &= \log \int \pi(\boldsymbol{\Lambda}, \boldsymbol{\theta}) \; d\boldsymbol{\theta} \\ &= \log \int q(\boldsymbol{\theta}) \frac{\pi(\boldsymbol{\Lambda}, \boldsymbol{\theta})}{q(\boldsymbol{\theta})} \; d\boldsymbol{\theta} \\ &\geq \int q(\boldsymbol{\theta}) \log \frac{\pi(\boldsymbol{\Lambda}, \boldsymbol{\theta})}{q(\boldsymbol{\theta})} \; d\boldsymbol{\theta} \quad \text{(Jensen's inequality)} \\ &= F(q, \boldsymbol{\theta}) \end{aligned} \tag{25}$$

It is obvious that this inequality becomes an equality when in place of the auxiliary distribution $q(\boldsymbol{\theta})$ the conditional posterior $\pi(\boldsymbol{\theta} \mid \boldsymbol{\Lambda}) = p(\boldsymbol{\theta} \mid \boldsymbol{\Lambda}, \mathbf{u}_Q, \mathcal{M})$ is selected. Given an estimate $\boldsymbol{\Lambda}^{(j)}$ at step j, this suggests iterating between an Expectation step (E-step) whereby we average with respect to $q^{(j)}(\boldsymbol{\theta}) = \pi(\boldsymbol{\theta} \mid \boldsymbol{\Lambda}^{(j)}, \mathbf{u}_Q, \mathcal{M})$ to evaluate the lower bound:

$$\begin{aligned} \textbf{E-step:} \quad F^{(j)}(q^{(j)}, \boldsymbol{\Lambda}) \quad &= \int q^{(s)}(\boldsymbol{\theta}) \log \pi(\boldsymbol{\Lambda}, \boldsymbol{\theta}) \; d\boldsymbol{\theta} \\ &- \int q^{(j)}(\boldsymbol{\theta}) \log q^{(j)}(\boldsymbol{\theta}) \; d\boldsymbol{\theta} \end{aligned} \tag{26}$$

and a Maximization step (M-step) with respect to $F^{(j)}(q^{(j)}, \boldsymbol{\Lambda})$ (and in particular the first part in Equation (26) since the second does no not depend on $\boldsymbol{\Lambda}$):

$$\begin{aligned} \textbf{M-step:} \quad \boldsymbol{\Lambda}^{(j+1)} \quad &= arg \max_{\boldsymbol{\Lambda}} F^{(j)}(q^{(j)}, \boldsymbol{\Lambda}) \\ &= arg \max_{\boldsymbol{\Theta}} E_{q^{(j)}(\boldsymbol{\theta})} [\log \pi(\boldsymbol{\Lambda}, \boldsymbol{\theta})] \\ &= arg \max_{\boldsymbol{\Lambda}} Q(\boldsymbol{\Lambda}^{(j)}, \boldsymbol{\Lambda}) \end{aligned} \tag{27}$$

Given the expression of the (unormalized) posterior in Equation (14), the aforementioned objective function $Q(\boldsymbol{\Lambda}^{(j)}, \boldsymbol{\Lambda})$ becomes:

$$\begin{aligned} Q(\boldsymbol{\Lambda}^{(j)}, \boldsymbol{\Lambda}) \quad &= E_{q^{(j)}(\boldsymbol{\theta})} [\log \pi(\boldsymbol{\Lambda}, \boldsymbol{\theta})] \\ &= E_{q^{(j)}(\boldsymbol{\theta})} [\log \textstyle\prod_{e=1}^{n_{el}} p(\mathbf{c}_e \mid \boldsymbol{\sigma}_e, \mathbf{u}_e, \lambda_e^2) \; p(\boldsymbol{\Lambda})] \\ &= E_{q^{(j)}(\boldsymbol{\theta})} [\textstyle\sum_{e=1}^{n_{el}} \log p(\mathbf{c}_e \mid \boldsymbol{\sigma}_e, \mathbf{u}_e, \lambda_e^2)] + E_{q^{(j)}(\boldsymbol{\theta})} [\log p(\boldsymbol{\Lambda})] \\ &= \textstyle\sum_{e=1}^{n_{el}} E_{q^{(j)}(\boldsymbol{\theta})} [\log p(\mathbf{c}_e \mid \boldsymbol{\sigma}_e, \mathbf{u}_e, \lambda_e^2)] + \log p(\boldsymbol{\Lambda}) \end{aligned} \tag{28}$$

While the second term in the expression above is essentially a penalty term arising from the prior on $\boldsymbol{\Lambda}$ (Equation (11)), the first term from Equation (10) leads to:

$$\begin{aligned} E_{q^{(j)}(\boldsymbol{\theta})} [\log p(\mathbf{c}_e \mid \boldsymbol{\sigma}_e, \mathbf{u}_e, \lambda_e^2)] = \quad &-\tfrac{n_\sigma}{2} \log \lambda_e^2 \\ &-\tfrac{1}{\lambda_e^2} E_{q^{(j)}(\boldsymbol{\theta})} [\| \boldsymbol{\sigma}_e - \mathbf{D}_e \mathbf{B}_e \mathbf{u}_e \|^2] \end{aligned} \tag{29}$$

It is evident that the M-step requires computation of the sufficient statistics Φ_e:

$$\Phi_e^{(j)} = E_{q^{(j)}(\boldsymbol{\theta})} \left[\| \boldsymbol{\sigma}_e - \mathbf{D}_e \mathbf{B}_e \mathbf{u}_e \|^2 \right] \tag{30}$$

i.e. the expected values (with respect to $q^{(j)}$) of the constitutive relation discrepancy in each of the elements $e = 1, \ldots, n_{el}$. Given the dependence amongst the components of $\boldsymbol{\Lambda}$ in the prior model, we propose an *incremental* version of the EM scheme ((Meng & Rubin 1993, Neal & Hinton 1998)) where rather than maximizing $Q(\boldsymbol{\Lambda}^{(j)}, \boldsymbol{\Lambda})$ in the M-step, we set $\boldsymbol{\Lambda}^{(j+1)}$ such that:

$$Q(\boldsymbol{\Lambda}^{(j)}, \boldsymbol{\Lambda}^{(j+1)}) \geq Q(\boldsymbol{\Lambda}^{(j)}, \boldsymbol{\Lambda}^{(j)}) \tag{31}$$

To that end we propose maximizing $Q(\boldsymbol{\Lambda}^{(j)}, \boldsymbol{\Lambda})$ with respect to a single component of $\boldsymbol{\Lambda}$ (i.e. $\lambda_e^2,\ e = 1, \ldots, n_{el}$) at a time while keeping the rest fixed. At each step, all the components of $\boldsymbol{\Lambda}$ were scanned and details on the computations entailed are provided in the Appendix.

The critical task is that of inference i.e. the calculation of the expectations with respect to $q^{(j)}(\boldsymbol{\theta})$ in the E-step (Equation (26) or Equation (29)). As mentioned earlier, the optimal choice for $q^{(j)}(\boldsymbol{\theta})$ is the (conditional) posterior $\pi(\boldsymbol{\theta} \mid \boldsymbol{\Lambda}^{(j)})$ which is analytically intractable as it can be readily be established from Equation (14). While suboptimal variational approximations can be employed (e.g. (Ghahramani & Attias 2000, Beal & Ghahramani 2003, Wainwright & Jordan 2008)), in this work we explore asymptotically exact, approximations based on MCMC sampling from the posterior (Robert & Casella 2004). If $\{\boldsymbol{\theta}^{(i,j)}\}_{i=1}^N$ denote N samples from such a Markov chain with the (conditional) posterior $q^{(j)}(\boldsymbol{\theta}) = \pi(\boldsymbol{\theta} \mid \boldsymbol{\Lambda}^{(j)})$ at iteration j as the target, then the E-step in Equation (26) can be substituted by:

$$Q(\boldsymbol{\Lambda}^{(j)}, \boldsymbol{\Lambda}) = \int q^{(j)}(\boldsymbol{\theta}) \log \pi(\boldsymbol{\Lambda}, \boldsymbol{\theta})\, d\boldsymbol{\theta} \approx \hat{Q}(\boldsymbol{\Lambda}^{(j)}, \boldsymbol{\Lambda}) = \frac{1}{N} \sum_{i=1}^{N} \log \pi(\boldsymbol{\Lambda}, \boldsymbol{\theta}^{(i,j)}) \tag{32}$$

The unanoivadable noise introduced in these estimates by MCMC might necessitate an exuberant number of samples N to obtain a robust algorithm particularly close to the maximum of $L(\boldsymbol{\Lambda})$ (Equation (24)). For that purpose we propose employing a stochastic approximation variant of the Robbins & Monro scheme (Robbins & Monro 1951, Cappé et al. 2005). Rather than increasing the simulation size N in order to reduce the variance, we compute a weighted average at the current and previous iterations. By employing a decreasing sequence of weights, information from the earlier iterations gets discarded gradually and more emphasis is placed on the recent iterations. As it is shown in (Delyon et al. 1999), this method converges with a fixed sample size N (even when $N = 1$). Convergence results that take into account the dependendence of the Markov chains at each EM-step have been obtained by constraining the sequence of $\boldsymbol{\Lambda}^{(j)}$ to some compact set $\mathcal{C}$ by means of a reprojection onto $\mathcal{C}$ (Kushner & Yin 2003). Even though this does not pose much problems in compuational practice, weakened conditions have been established in (Andrieu et al. 2005, Liang et al. 2007). In particular, rather than

using $\hat{Q}(\mathbf{\Lambda}^{(j)}, \mathbf{\Lambda})$ (which according to Equation (32) approximates $Q(\mathbf{\Lambda}^{(j)}, \mathbf{\Lambda})$) in the M-step (Equation (27)), we use:

$$\tilde{Q}(\mathbf{\Lambda}^{(j)}, \mathbf{\Lambda}) = (1-\gamma_j)\tilde{Q}(\mathbf{\Lambda}^{(j-1)}, \mathbf{\Lambda}) + \gamma_j \hat{Q}(\mathbf{\Lambda}^{(j)}, \mathbf{\Lambda}) \tag{33}$$

where the sequence of weights $\{\gamma_j\}$ is such that $\sum_{j=1}^{\infty} \gamma_j = \infty$ and $\sum_{j=1}^{\infty} \gamma_j^2 < \infty$ ¶. As it can be seen from Equations (28), (29) and (30) in order to estimate the weighted average in Equation (33), it suffices to keep track of the weighted averages $\tilde{\Phi}_e^{(j)}$ of the sufficient statistics $\Phi_e^{(j)}$ (Equation (30)):

$$\tilde{\Phi}_e^{(j)} = (1-\gamma_j)\tilde{\Phi}_e^{(j-1)} + \gamma_j \Phi_e^{(j)} \tag{34}$$

The MCMC steps can be carried out using Gibbs sampling with respect to each of the components of $\boldsymbol{\theta}$ i.e. ν^2 $\mathbf{u}$, $\boldsymbol{\sigma}$ and $\{\mathbf{D}_e\}_{e=1}^{n_{el}}$ which require the conditional distributions enumerated in the previous subsection (i.e. Equations (16), (17), (20) and (22)). It is worth pointing out that the $n \times n$ system of linear equations *does not need to be solved* (which has a cost of $O(n^3)$ operations) at any stage as in traditional inverse problems. If J is the total number of EM iterations and N is the number of MCMC steps at each iteration, then sampling from the aforementioned conditionals implies:

- the inversion and Cholesky factorization of $\mathbf{C}_u$ in order to generate samples of $\mathbf{u}$. This must be repeated at *every* MCMC step since $\{\mathbf{D}_e\}$ are updated. The cost of this operation is $O(J\ N\ n^3)$.
- the Cholesky factorization of $\mathbf{C}_\sigma$ in order to generate samples of $\boldsymbol{\sigma}$. This must be repeated at *every* EM iteration and *not at every* MCMC step since $\mathbf{C}_\sigma$ solely depends on $\mathbf{\Lambda}$. The cost of this operation is $O(J\ (n_{el} n_\sigma)^3)$ where where n_σ is the number of stress components ($n_\sigma = 6$ in three dimensions, $n_\sigma = 3$ in plane stress/strain etc).

In order to reduce the cost associated with these operations one can employ block-Gibbs updates with respect to each of the components of $\mathbf{u}$ (or blocks of $\mathbf{u}$) rather than updating the whole vector at once. As it is demonstrated in the sequence the cost of such a scheme is $O(J\ N\ n(n_{el} n_\sigma))$. The mixing is obviously slower than the full updates and as a consequence the variance in the MCMC estimates is larger. In general therefore more EM iterations (assuming the same number of samples N are used at each iteration) are needed to converge. Nevertheless the linear scaling with J constitutes such a scheme more efficient. Similar block-Gibbs updates can be carried out for $\boldsymbol{\sigma}$ reducing the cost associated with this task to $J\ N\ n(n_{el} n_\sigma))$. The conditional posteriors for performing block-Gibbs moves are described in the sequence.

Let $\begin{bmatrix} u_i \\ \mathbf{u}_{-i} \end{bmatrix}$ be a partitioning of $\mathbf{u}$ with respect to component i^+. Let also $\mathbf{Q} = [\mathbf{Q}_i \mid \mathbf{Q}_{-i}]$, $\mathbf{C} = [\mathbf{C}_i \mid \mathbf{C}_{-i}]$ the corresponding partitioning of the matrices appearing

¶ A family of such sequences that was used in this work is $\gamma_j = \frac{1}{j^p}$ with $1/2 < p \leq 1$. the value of $p = 0.51$ was employed

$^+$ An identical procedure can be followed when u_i corresponds to a block of $\mathbf{u}$

in Equation (3) and Equation (19). Then the conditional posterior of u_i from Equation (14) is:

$$u_i \mid \mathbf{u}_{-i}, \nu^2, \boldsymbol{\sigma}, \{\mathbf{D}_e, \lambda_e^2\}_{e=1}^{n_{el}} \sim \mathcal{N}(\mu_{u_i}, \sigma_{u_i}^2) \tag{35}$$

where:

$$\sigma_{u_i}^{-2} = \mathbf{C}_i^T \boldsymbol{\Lambda}^{-1} \mathbf{C}_i + \frac{1}{\nu^2} \mathbf{Q}_i^T \mathbf{Q}_i \tag{36}$$

$$\mu_{u_i} = \sigma_{u_i}^2 \left(\mathbf{C}_i^T \boldsymbol{\Lambda}^{-1} (\boldsymbol{\sigma} - \mathbf{C}_{-i} \mathbf{u}_{-i}) + \frac{1}{\nu^2} \mathbf{Q}_i^T (\mathbf{u}_Q - \mathbf{Q}_{-i} \mathbf{u}_{-i}) \right) \tag{37}$$

It is noted that the leading order of computational operations for updating succesively all components of $\mathbf{u}$ as above is $O(n(n_{el} n_\sigma))$. This is approximately one order less than the $O(n^3)$ cost associated with the full update (Equation (17)), considering that the dimension of the stress vector $n_{el} n_\sigma$ is comparable to n.

3. Numerical examples

In this section we report results on the accuracy and performance of the algorithm on two-dimensional elastography problems on synthetic data obtained for the configuration depicted in Figure 2 (Albocher et al. 2009, Barbone et al. 2010) where the boundary displacements normal to the walls are prescribed . We intend to provide a clinical validation of the approach in a future study.

We assume an isotropic elastic material with Poisson's ratio $\nu = 0.5$ (incompressible) and employ the selective reduced integration quadrilateral elements for the forward problem (Hughes 1980, Hughes 2000). We examine two distributions for the elastic modulus occurring in elliptic and circular inclusions. In the first problem (Figure 3) the emphasis is on demonstrating the capabilities of the proposed method in identifying the ground truth as well as providing probabilistic confidence metrics particularly in the presence of noise. In the second case (Figure 7) the emphasis is on detecting and quantifying model discrepancies in the sense described in section 2. It is noted that in all cases apart from the identification of material properties, a direct output of the computations is the stress distribution. It is finally noted that in order to generate the displacement data, the forward problem was solved with a randomly generated mesh consisting of $10,000$ elements.

The following values were used for the parameters appearing in prior models described previously:

- $\sigma_z^2 = 100$ (Equation (11)) which corresponds to a diffuse prior and $\sigma_u^2 = 1$ in Equation (15). The latter was selected based on the magnitude of the prescribed boundary displacements in Figure 2.
- $d_0 = 0.1$ for the correlation-length parameter appearing in the $\mathbf{H}$ (Equation (12)) and $\mathbf{V}$ (Equation (15)). Numerical evidence suggested that the effect of this parameter was minimal when varied in the range $[0.01, 0.5]$ given that the characteristic dimension of the problem domain is 1.

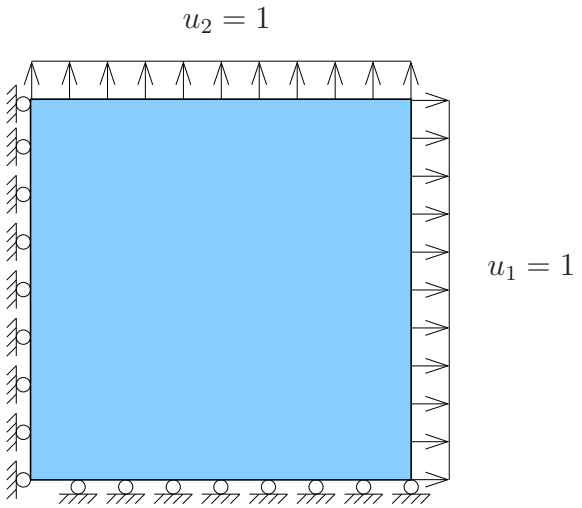

Figure 2. Problem configuration used in both examples 1 and 2 (Albocher et al. 2009)

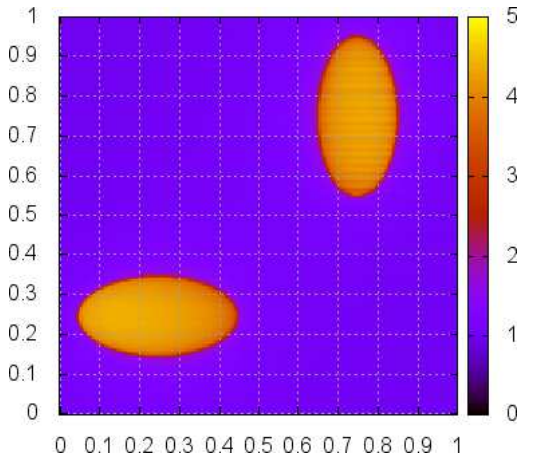

Figure 3. Example 1 - Elastic modulus E spatial distribution: In the inclusions $E = 5$ whereas in the rest of the domain $E = 1$

- an uniformative Jeffry's prior was adopted for the observation noise variance ν^2 (Equation (16)) with $\alpha_\nu = 2$ and $\beta_\nu = 0$.

With regards to the EM scheme, at each iteration $N = 10$ MCMC updates of all model parameters were performed and iterations were terminated when the *relative increase* in the objective $Q(\mathbf{\Lambda}^{(j)}, \mathbf{\Lambda})$ in Equation (28) was less or equal than $\epsilon = 0.001$, i.e. $\frac{|Q(\mathbf{\Lambda}^{(j)},\mathbf{\Lambda}) - Q(\mathbf{\Lambda}^{(j-1)},\mathbf{\Lambda})|}{|Q(\mathbf{\Lambda}^{(j-1)},\mathbf{\Lambda})|} < \epsilon$.

3.1. Example 1

The first scenario involves two elliptical inclusions centered at $(0.25, 0.25)$ and $(0.75, 0.75)$ with principal axes 0.1 and 0.2. with a contrast ratio 5 : 1 in the elastic modulus (Figure 3). A useful outcome of the numerical investigations was the fact that the overall inference and learning process can be greatly accelerated by operating on a sequence of discretizations with increased refinement. In particular, initially a coarse mesh is adopted with few nodes and elements where the proposed EM scheme is applied. The parameter values learned (i.e. $\mathbf{\Lambda}$) are used as the initial values for a refined mesh. The MCMC chains with respect to the other model parameters at the new mesh are initiated from samples drawn at the coarser mesh. It was found that this led to a reduction of the number of EM iterations needed to achieve convergence and significant acceleration since the order of operations at coarse meshes is smaller. For that purpose, we report in this problem the results obtained at three different resolutions employing a regular mesh with 5×5 , 10×10 and 20×20 elements. A potentially important implication involves the possibility of *adaptive refinement* where the mesh can be refined at selected regions of the problem domain where further information is needed as determined by the inferences produced at coarser resolutions (Arnold et al. 2010).

Figure 4 depicts the posterior mean as well as the posterior quantiles at 5% and 95% for the elastic modulus at these three resolutions and in the absence of noise in the data. It is readily observed that the proposed scheme can identify the ground truth as well as provide posterior credible intervals on the inferences made. These are more clearly depicted in Figure 6(a) which presents the results along the diagonal from $(0, 0)$ to $(1, 1)$.

We also investigated the performance of the algorithm in the presence of zero mean, Gaussian noise and in particular with a Signal-to-Noise-Ratio $SNR = 40dB$ which is typical for ultrasound systems (Doyley et al. 2006, Oberai et al. 2004). The results are shown in Figure 5 in terms of posterior mean and posterior quantiles. As it can also be seen in Figure 6(b), the algorithm is able to quantify the uncertainty introduced by the presence of noise and posterior bounds provided enclose the ground truth. Finally Figure **??** depicts randomly selected samples drawn at various iterations of the EM scheme (for the finest resolution 20×20) that demonstrate the evolution of the learning algorithm proposed.

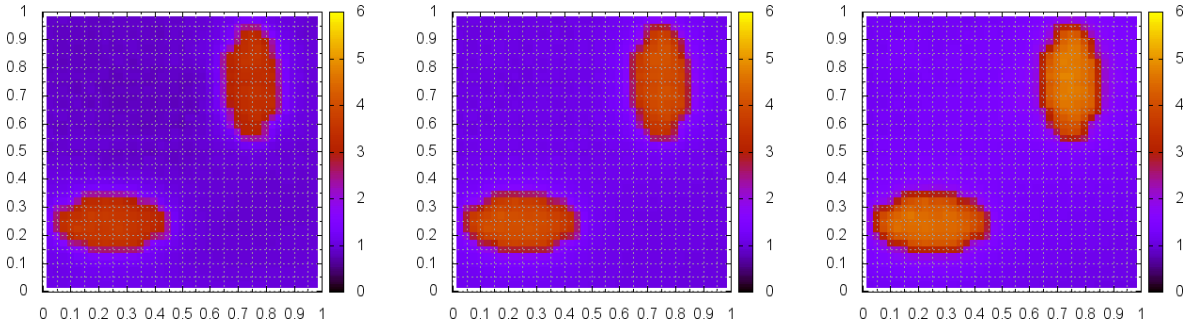

(a) 20x20: 5% quantile (b) 20x20: posterior mean (c) 20x20: 95% quantile

Figure 4. Example 1: Posterior statistics of the elastic modulus distribution for noiseless data.

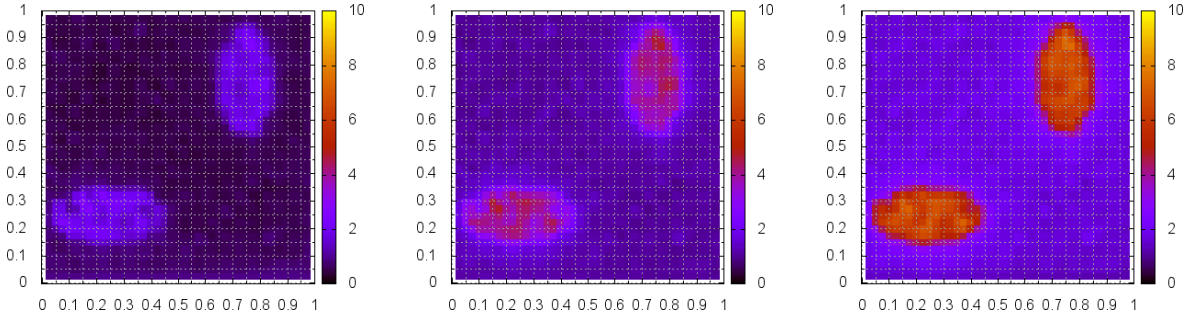

(a) 20x20: 5% quantile (b) 20x20: posterior mean (c) 20x20: 95% quantile

Figure 5. Example 1: Posterior statistics of the elastic modulus distribution for noisy data with SNR=$40dB$.

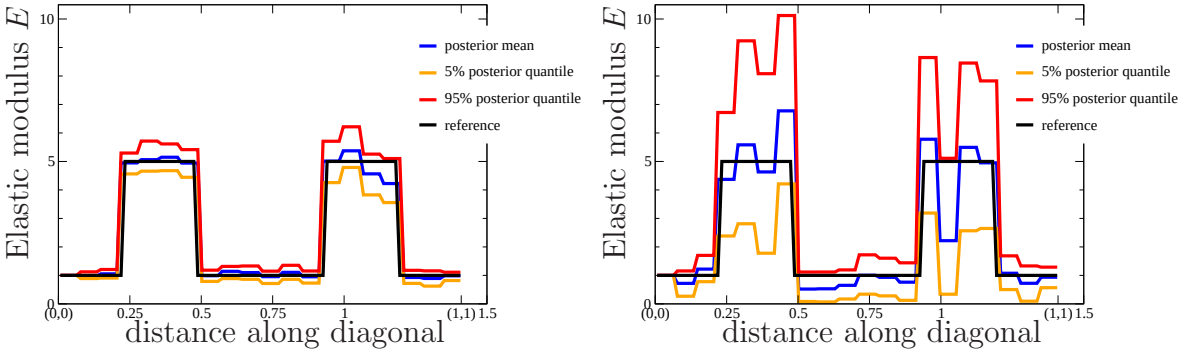

(a) no noise (b) SNR=40

Figure 6. Example 1: Posterior statistics of the elastic modulus distribution along the diagonal from $(0, 0)$ to $(1, 1)$.

3.2. Example 2

The primary goal in the second example is to demonstrate the capability of quantifying model discrepancy in the constitutive equation. In particular we consider the synthetic data generated by the material distribution in Figure 7 *. The circular inclusion centered at $(0.5, 0.5)$ with radius 0.2 is assumed to have an elastic modulus that is 5 times larger than the rest of domain. We further assumed a square region on the top left corner $[0, 0.2] \times [0.8, 1]$ where rather than an *isotropic*, elastic material we employed an anisotropic constitutive matrix $\mathbf{D} = \begin{bmatrix} 10 & -5 & -5 \\ -5, & 20. & -5 \\ -5 & -5 & 100 \end{bmatrix}$. While this is a valid constitutive model (i.e. $\mathbf{D}$ is positive definite) it is obviously inconsistent with the isotropic assumption made in the model used to identify material properties. While other inversion schemes might be able to find an elastic modulus corresponding to an isotropic material that fits adequately the observed displacemnts, they would be unable to identify that the model employed is inadequate. As a result erroneous conclusions would be drawn about the state of the material at this portion of the problem domain.

Figure 8 depicts the learned values of the the parameters $\mathbf{\Lambda} = \{\lambda_e^2\}$ (Equation (10)) which express the magnitude of model error over each element of the domain. Both in the absence of noise and when SNR=$40dB$, the algorithm clearly identifies a significant model error in the region on the top-left corner. It is noted that that the λ_e^2 values in this region are 2 to 4 orders of magnitude larger than in the rest of the problem domain. Despite the model inadequacy the algorithm correctly identifies the presence of the circular inclusion as it can be seen in Figure 9 and more clearly in Figure

* In the circular inclusion $E = 5$, in the subdomain $[0, 0.2] \times [0.8, 1]$ we employed a constitutive matrix $\mathbf{D} = \begin{bmatrix} 10 & -5 & -5 \\ -5, & 10. & -5 \\ -5 & -5 & 100 \end{bmatrix}$ whereas in the rest of the domain $E = 1$

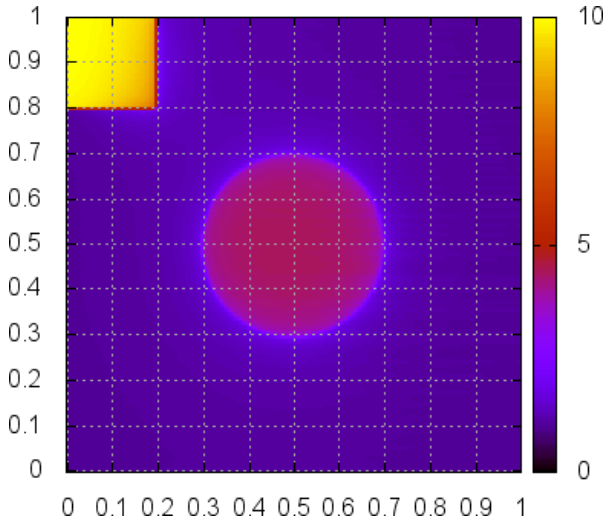

Figure 7. Example 2 - Elastic modulus E spatial distribution

10 which shows the elastic modulus variation along the diagonal from $(0, 1)$ to $(1, 0)$. It is particularly interesting to note that even though the isotropic elastic constitutive model endowed in the inversion scheme is inadequate at least for a subdomain of the problem, the proposed scheme can correctly identify the stresses (pressure and shear) in the whole domain as it can be seen in Figures 11 and 12. These depict the ground truth in comparison with the posterior means obtained with no noise and for SNR=40dB. The posterior quantiles (which are omitted herein for economy of space) fully envelop the ground truth.

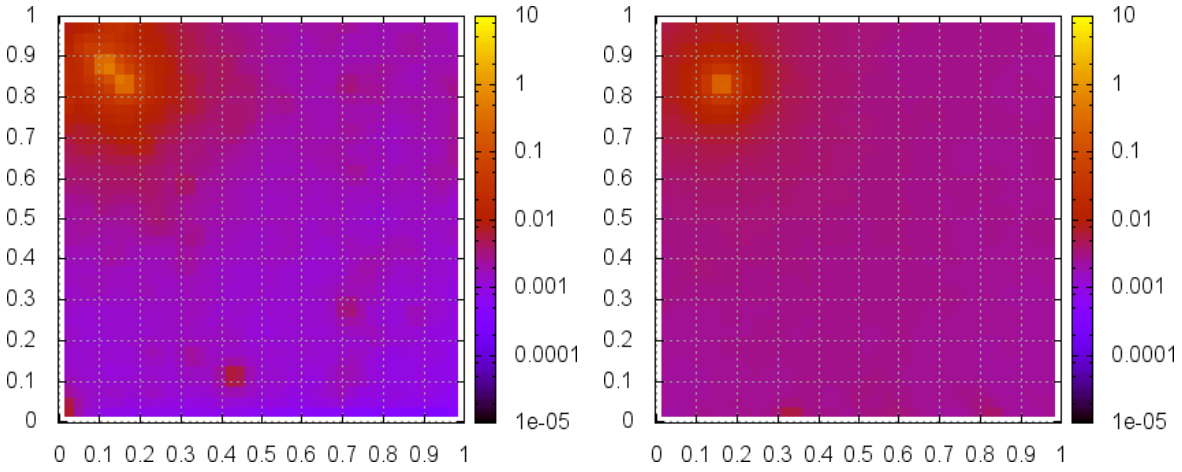

(a) no noise (b) SNR=40

Figure 8. Example 2: Model discrepancies/errors $\{\lambda_e^2\}_{e=1}^{n_{el}}$ for a) no noise, and b) SNR=40dB (in log-scale)

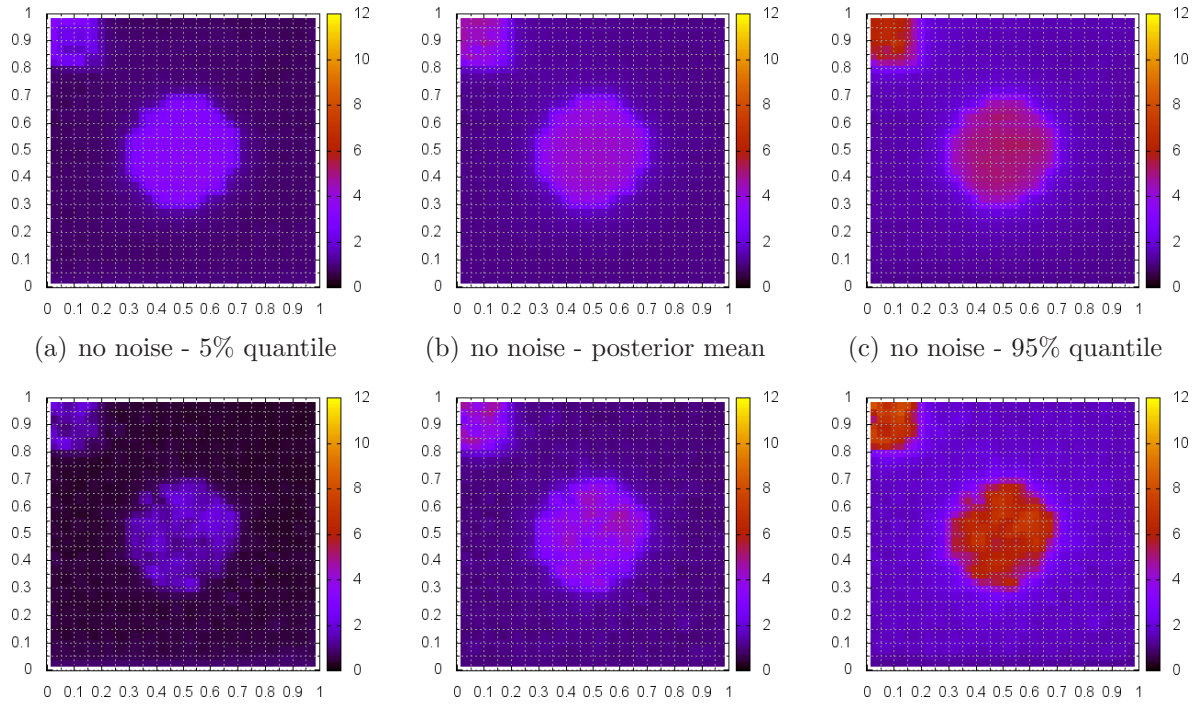

(a) no noise - 5% quantile (b) no noise - posterior mean (c) no noise - 95% quantile

(d) SNR=40 - 5% quantile (e) SNR=40 - posterior mean (f) SNR=40 - 95% quantile

Figure 9. Example 2: Posterior statistics of the elastic modulus distribution when

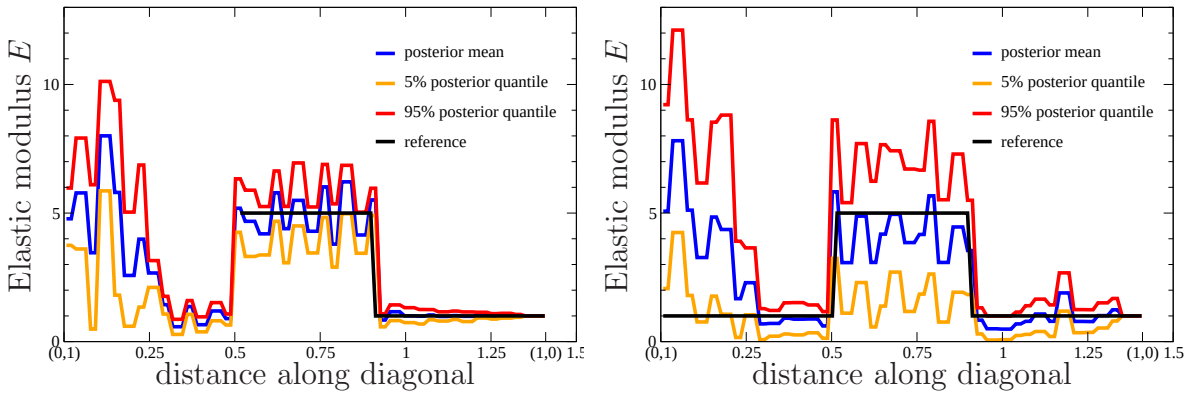

(a) no noise (b) SNR=40

Figure 10. Example 2: Posterior statistics of the elastic modulus distribution along the diagonal from $(0, 1)$ to $(1, 0)$

4. Conclusions

While existing stochastic (Bayesian) strategies for the solution of inverse problems associated with the identification of material properties in biomechanics are able to account for various sources of uncertainty in the problem, they are generally deficient in terms of assessing model fidelity. We proposed an intrusive formulation that incorporates the various model equations in the likelihood (posterior) and is capable of inferring model

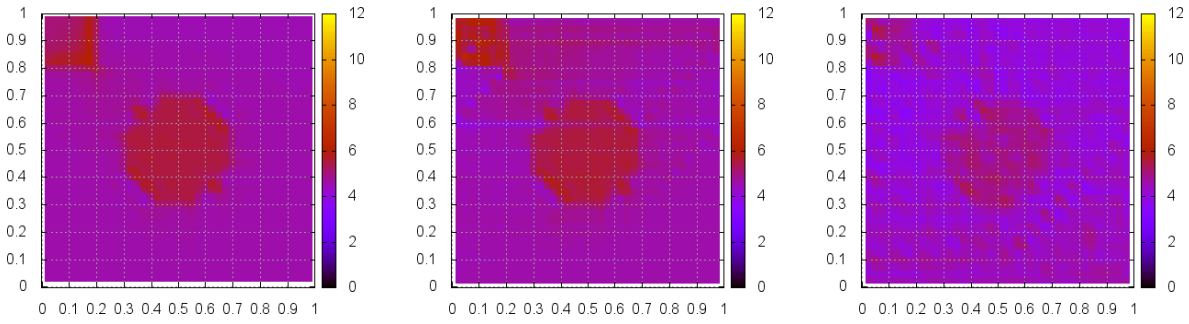

(a) reference (ground truth) (b) no noise - posterior mean (c) SNR=40 - posterior mean

Figure 11. Example 2: Comparison of pressure's spatial distribution with the posterior means obtained when data have no noise and for SNR=40dB.

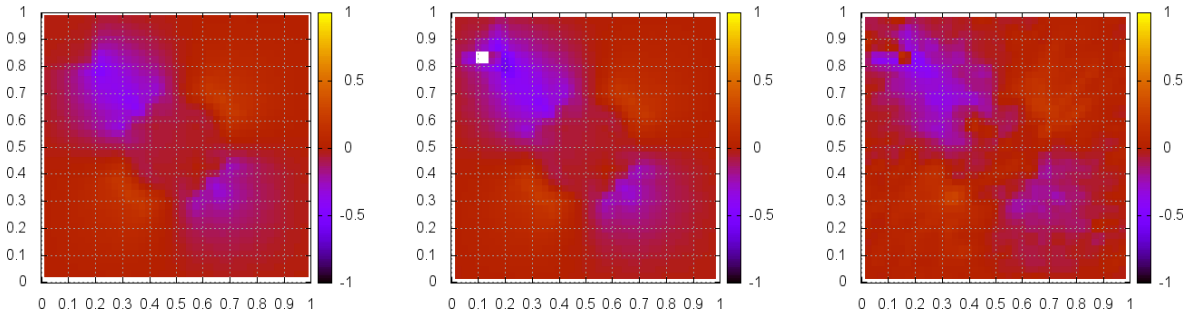

(a) reference (ground truth) (b) no noise - posterior mean (c) SNR=40 - posterior mean

Figure 12. Example 2: Comparison of shear stress' σ_{xy} spatial distribution with the posterior means obtained when data have no noise and for SNR=40dB.

discrepancies from noisy displacement data. In contrast to direct methods, it does not require imputations of strains nor their derivatives. It provides probabilistic confidence metrics (credible intervals) that can be very useful to the analyst as well as probabilistic estimates of the (unobserved) stresses/pressures. We discussed a scalable computational framework which can be greatly accelerated by employing a multi-resolution strategy. The latter could be utilized in order to propose adaptively, refinements of the discretized domain which we intend to explore in the future. Current investigations also involve extending this approach to dynamic settings where the parameter vector should include velocities and accelerations in addition to displacements, and the model equations should include the time-integration scheme adopted.

Appendix: Maximization with respect to $\mathbf{\Lambda}$

This section describes the computations involved during the Maximization step of the EM algorithm described in section 2. In particular according to Equations (28), (29), (30) and the prior model in Equation (11), this entails a maximization with respect to $\mathbf{\Lambda} = \{\lambda_e^2\}_e$ of:

$$\begin{aligned} Q(\mathbf{\Lambda}^{(j)}, \mathbf{\Lambda}) &= -\tfrac{n_\sigma}{2} \sum_{e=1}^{n_{el}} \log \lambda_e^2 - \tfrac{1}{\lambda_e^2} \sum_{e=1}^{n_{el}} \Phi_e^{(j)} + \log p(\mathbf{\Lambda}) \\ &= -\tfrac{n_\sigma}{2} \sum_{e=1}^{n_{el}} \log \lambda_e^2 - \tfrac{1}{\lambda_e^2} \sum_{e=1}^{n_{el}} \Phi_e^{(j)} - \tfrac{1}{2} \mathbf{Z}^T \mathbf{W} \mathbf{Z} \end{aligned} \tag{38}$$

It is reminded that the vector $\mathbf{Z} = \{z_e\}_{e=1}^{n_{el}}$ contains the log values of $\mathbf{\Lambda}$ i.e. $z_e = \log \lambda_e^2$. Rather than solving an optimization in the n_{el}-dimensional space at each iteration j, we perform successive updates of each λ_e^2 or z_e while keeping the remaining fixed. This *incremental* version of the EM algorithm entails performing n_{el} optimizations of one-dimensional functions. We propose carrying out the latter task with respect to z_e (as they are allowed to take any value on the real axis in contrast to λ_e^2 which must be positive) and employ a standard Newton-Raphson scheme. This requires the first and second order derivatives of the objective function above which are given by:

$$\frac{\partial Q(\mathbf{\Lambda}^{(j)}, \mathbf{\Lambda})}{\partial z_e} = -\frac{3}{2} + \frac{\Phi_e^{(j)}}{2} e^{-z_e} - \frac{z_e - \mu_{z_e}}{\sigma_{z_e}^2} \tag{39}$$

and:

$$\frac{\partial^2 Q(\mathbf{\Lambda}^{(j)}, \mathbf{\Lambda})}{\partial z_e^2} = -\frac{\Phi_e^{(j)}}{2} e^{-z_e} - \frac{1}{\sigma_{z_e}^2} \tag{40}$$

where:

$$\begin{aligned} \sigma_{z_e}^2 &= 1/W_{e,e} \\ \mu_{z_e} &= -\tfrac{1}{W_{e,e}} \sum_{k \neq e} W_{e,k} z_k \end{aligned} \tag{41}$$

It can be easily seen that the second derivative is always, strictly negative $\frac{\partial^2 Q(\mathbf{\Lambda}^{(j)}, \mathbf{\Lambda})}{\partial z_e^2} < 0$ and therefore the problem is convex.

References

Albocher U, Oberai A, Barbone P & Harai I 2009 *Comput. Methods Appl. Mech. Engrg.* **198**, 24122420.

Andrieu C, Moulines E & Priouret P 2005 *SIAM J. Control Optim.* .

Arnold A, Reichling S, Bruhns O T & Mosler J 2010 *PHYSICS IN MEDICINE AND BIOLOGY* **55**(7), 2035–2056.

Bamber J, Barbone P, Bush N, Cosgrove D, Doyely M, Fuechsel F, Meaney P, Miller N, Shiina T & Tranquart F 2002 *IEICE TRANSACTIONS ON INFORMATION AND SYSTEMS* **E85D**(1), 5 – 14.

Banerjee B, Roy D & Vasu R M 2009 *PHYSICS IN MEDICINE AND BIOLOGY* **54**(2), 285–305.

Barbone P E, Rivas C E, Harari I, Albocher U, Oberai A A & Zhang Y 2010 *INTERNATIONAL JOURNAL FOR NUMERICAL METHODS IN ENGINEERING* **81**(13), 1713–1736.

Beal M J & Ghahramani Z 2003 *Bayesian Statistics* (7).

Besag J & Green P 1993 *J. Royal Statist. Soc. Ser. B, Methodological* **55**, 25–37.

Besag J & Kooperberg C 1995 *BIOMETRIKA* **82**(4), 733–746.

Besag J, York J & Mollié 1991 *Annals of the Institute of Statistical Mathematics* **43**, 1–59.

Bonnet M & Constantinescu A 2005 *Inverse Problems* **21**, R1–R50.

Burnside E, Hall T, Sommer A, Hesley G, Sisney G, Svensson W, Fine J, Jiang J & Hangiandreou N 2007 **245**(2), 401 – 410.

Cappé O, Moulines E & Rydén T 2005 *Inference in Hidden Markov Models* Springer-Verlag.

Delyon B, Lavielle M & Moulines E 1999 *The Annals of Statistics* **27**, 94–128.

Dempster A, Laird N & Rubin D 1977 *J. Roy. Statist. Soc. Ser. B* **39**(1), 1–38.

Deraemaeker A, Ladéveze P & Romeuf T 2004 *Eng. Comput.* pp. 21808–33.

Doyley M, Srinivasan S, Dimidenko E, Soni N & Ophir J 2006 *PHYSICS IN MEDICINE AND BIOLOGY* **51**(1), 95–112.

Fehrenbach J, Masmoudi M, Souchon R & Trompette P 2006 *INVERSE PROBLEMS* **22**(3), 1055–1069.

Feissel P & Allix O 2007 *Comput. Methods Appl. Mech. Eng* pp. 1968–83.

Ganne-Carrié N, Ziol M, de Ledinghen V, Douvin C, Marcellin P, Castera L, Dhumeaux D, Trinchet J & Beaugrand M 2006 *Hepatology* **44**(6), 15111517.

Garra B, Cespedes E, Ophir J, Spratt S, Zuurbier R, Magnant C & Pennanen M 1997 **202**(1), 79 – 86.

Ghahramani Z 2001 *Journal of Pattern Recognition and Artificial Intelligence* **15**(1), 9–42.

Ghahramani Z & Attias H 2000 Online variational bayesian learning. Slides from talk presented at NIPS 2000 workshop on Online Learning.

Giuseppetti G, Martegani A, Di cioccio B & Baldassarre S 2005 *RADIOLOGIA MEDICA* **110**(1-2), 69 – 76.

Gockenbach M S, Jadamba B & Khan A A 2008 *INVERSE PROBLEMS IN SCIENCE AND ENGINEERING* **16**(3), 349–367.

Gockenbach M S & Khan A A 2005 *JOURNAL OF INDUSTRIAL AND MANAGEMENT OPTIMIZATION* **1**(4), 487–497.

Gokhale N H, Barbone P E & Oberai A A 2008 *INVERSE PROBLEMS* **24**(4).

Hall T, Zhu Y & Spalding C 2003 *ULTRASOUND IN MEDICINE AND BIOLOGY* **29**(3), 427 – 435.

Higdon D 1997 *Journal of the American Statistical Association* .

Higdon D, Nakhleb B, Gattiker J & B. W 2008 *COMPUTER METHODS IN APPLIED MECHANICS AND ENGINEERING* **197**(29-32), 2431 – 2441.

Hughes T J R 1980 *Int. J. Numer. Methods Eng.* **15**, 1413.

Hughes T J R 2000 *The Finite Element MethodLinear Static and Dynamic Finite Element Analysis* Dover.

Itoh A, Ueno E, Tohno E, Kamma H, Takahashi H, Shiina T, Yamakawa M & Matsumura T 2006 **239**(2), 341 – 350.

Kaipio J & Somersalo E 2005 *Computational and Statistical Methods for Inverse Problems* Springer-Verlag, New York.

Kennedy M & O'Hagan A 2001 *JOURNAL OF THE ROYAL STATISTICAL SOCIETY SERIES B-STATISTICAL METHODOLOGY* **63**, 425 – 450.

Kohn R & Vogelius M 1984 *Pure Appl. Math.* **37289**.

Kunsch H 1987 *BIOMETRIKA* **74**, 517.

Kushner H & Yin G 2003 *Stochastic approximation and recurcive algorithms and applications* Springer.

Ladéveze P & Chouaki A 1999 *Inverse Problems* **15**, 49–58.

Liang F, Liu C & Carroll R J 2007 *J. Amer. Statist. Assoc.* **102**, 305–320.

Liew H & Pinsky P 2005 *FINITE ELEMENTS IN ANALYSIS AND DESIGN* **41**(7-8), 778–799.

Liu Y, Sun L & Wang G 2005 *IEEE TRANSACTIONS ON MEDICAL IMAGING* **24**(10), 1323–1333.

McLaughlin J & Renzi D 2006 *INVERSE PROBLEMS* **22**(2), 707–725.

Meng X L & Rubin D B 1993 *Biometrika* **80**, 267–278.

Neal R & Hinton G E 1998 *in* 'Learning in Graphical Models' Kluwer Academic Publishers pp. 355–368.

Oberai A A, Gokhale N H, Goenezen S, Barbone P E, Hall T J, Sommer A M & Jiang J 2009 *PHYSICS IN MEDICINE AND BIOLOGY* **54**(5), 1191–1207.

Oberai A, Gokhale N, Doyley M & Bamber J 2004 *PHYSICS IN MEDICINE AND BIOLOGY* **49**(13), 2955–2974.

Olson L G & Throne R D 2010 *INVERSE PROBLEMS IN SCIENCE AND ENGINEERING* **18**(6), 813–834.

Ophir J, Cespedes I, Ponnekanti H, Yazdi Y & Li X 1991 *ULTRASONIC IMAGING* **13**(2), 111 – 134.

Park E & Maniatty A M 2009 *INVERSE PROBLEMS IN SCIENCE AND ENGINEERING* **17**(5), 605–626.

Parker K, Doyley M M & Rubens D 2011 *Phys. Med. Biol.* .

Regner D, Hesley G, Hangiandreou N, Morton M, Nordland M, Meixner D, Hall T, Farrell M, Mandrekar J, Harmsen W & Charboneau J 2006 **238**(2), 425 – 437.

Robbins H & Monro S 1951 *The Annals of Mathematical Statistics* **22**, 400407.

Robert C P & Casella G 2004 *Monte Carlo Statistical Methods* 2nd edn Springer New York.

Thomas A, Fischer T, Frey H, Ohlinger R, Grunwald S, Blohmer J, Winzer K, Weber S, Kristiansen G, Ebert B & Kummel S 2006 *ULTRASOUND IN OBSTETRICS AND GYNECOLOGY* **28**(3), 335 – 340.

Wainwright M & Jordan M 2008 *in* 'Foundations and Trends in Machine Learning' Vol. 1 of *1-305* pp. 1–305.

Zhang Y, Hall L, Goldgof D & Sarkar S 2006 *IEEE TRANSACTIONS ON EVOLUTIONARY COMPUTATION* **10**(3), 341–357.

Zhi H, Ou B, Luo B, Feng X, Wen Y & Yang H 2007 *JOURNAL OF ULTRASOUND IN MEDICINE* **26**(6), 807 – 815.